

PATENT

TRADEMARK

COPYRIGHT

TRADE SECRETS

UNFAIR COMPETITION

LICENSING

COMPLEX LITIGATION

MAILING ADDRESS:

P.O. BOX 45862

SALT LAKE CITY, UT 84145

INTERNET

HOME PAGE: <http://www.wnspat.com>GENERAL E-MAIL: info@wnspat.com

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Box: PATENT APPLICATION
Assistant Commissioner for Patents
Washington, DC 20231

TRANSMITTAL FOR PATENT APPLICATION

Sir:

Transmitted herewith for filing under 37 C.F.R. § 1.53(b) is a United States patent application entitled ARBITRATING AND SERVICING POLYCHRONOUS DATA REQUESTS IN DIRECT MEMORY ACCESS in the name of the following inventor(s):

Inventor: Donald M. Gray, III
Citizenship: United States
Address: 926 Sanchez Street
San Francisco, CA 94114

Inventor: Zaigham Ahsan
Citizenship:
Address:

Enclosed are the following:

☒ A specification, claims, abstract, and cover page in total comprising thirty-six (36) pages.

☒ Four (4) sheets of drawings.

☒ A Certificate of Express Mailing by "Express Mail" certifying a filing date by use of Express Mail Label No. EL565973450US.

X No fees are enclosed.

X No Declaration, Power of Attorney, and Petition is enclosed.

____ Priority to _____ (country) Application Serial No. _____ filed on
_____ is claimed under 35 U.S.C. § 119.

____ A certified copy of foreign priority Application Serial No. _____.

____ An Associate Power of Attorney.

Please address all future correspondence in connection with the above-identified patent application to the attention of the undersigned.

Dated this 31st day of July 2000.

Respectfully submitted,



CARL T. REED
Attorney for Applicant
Registration No. 45,454



022913

PATENT TRADEMARK OFFICE

Customer No.

CTR:dfw

G:\DATA\WPDOCSRN\WEBTV\OTHERDOC\0731patent_trx_73.doc

007E20" E47B2360

"Express Mail" Mailing Label No.: EL565973450US

Date of Deposit: July 31, 2000

CERTIFICATE OF MAILING BY "EXPRESS MAIL"

- Patent Application in the names of Donald M. Gray, III and Zaigham Ahsan for ARBITRATING AND SERVICING POLYCHRONOUS DATA REQUESTS IN DIRECT MEMORY ACCESS (36 pages)
- Drawings (4 sheets)
- Transmittal Letter (2 pages)
- Postcard

Dated this 31st day of July 2000.

Respectfully submitted,

Carl Reed

CARL T. REED
Attorney for Applicant
Registration No. 45,454

WORKMAN, NYDEGGER & SEELEY
1000 Eagle Gate Tower
60 East South Temple
Salt Lake City, Utah 84111
Telephone: (801) 533-9800

CTR:dfw

G:\DATA\WPDOCSRN\WEBTV\OTHERDOC\0731patent trx 73.doc

Variable	Mean	SD	Min	Max
Age	38.5	12.5	18	65
Gender	0.5	0.5	0	1
Marital status	0.5	0.5	0	1
Education	12.5	2.5	9	16
Income	3500	1500	1000	8000
Health status	0.5	0.5	0	1
Smoking status	0.3	0.5	0	1
Alcohol consumption	0.2	0.4	0	1
Exercise frequency	0.5	0.5	0	1
Stress level	4.5	1.5	1	7
Sleep quality	0.5	0.5	0	1
Work satisfaction	0.5	0.5	0	1
Life satisfaction	0.5	0.5	0	1
Depression score	10.5	5.5	0	30
Anxiety score	12.5	6.5	0	35
Quality of life score	75.5	15.5	30	100

UNITED STATES PATENT APPLICATION

of

Donald M. Gray, III

and

Zaigham Ahsan

for

**ARBITRATING AND SERVICING POLYCHRONOUS DATA
REQUESTS IN DIRECT MEMORY ACCESS**

COFFEE

WORKMAN, NYDEGGER & SEELEY
A PROFESSIONAL CORPORATION
ATTORNEYS AT LAW
1000 EAGLE GATE TOWER
60 EAST SOUTH TEMPLE
SALT LAKE CITY, UTAH 84111

BACKGROUND OF THE INVENTION

1. The Field of the Invention

The present invention relates to systems and methods for transferring data to and from memory in a computer system. More particularly, the present invention relates to systems and methods for servicing the data and memory requirements of system devices by arbitrating the data requests of those devices.

2. The Prior State of the Art

An important operational aspect of a computer or of a computer system is the need to transfer data to and from the memory of the computer. However, if the computer's processor is used to perform the task of transferring data to and from the computer's memory, then the processor is unable to perform other functions. When a computer is supporting high speed devices that have significant memory needs, the processor bears a heavy load if the processor is required to copy data word by word to and from the computer's memory system for those devices. As a result, using the processor to transfer data in this manner can consume precious processing time.

A solution to this problem is Direct Memory Access (DMA). A DMA controller essentially relieves the processor of having to transfer data to and from memory by permitting a device to transfer data to or from the computer's memory without the use of the computer's processor. A significant advantage of DMA is that large amounts of data may be transferred before generating an interrupt to the computer to signal that the task is completed. Because the DMA controller is transferring data, the processor is therefore free to perform other tasks.

1 As computer systems become more sophisticated, however, it is becoming
2 increasingly evident that there is a fundamental problem between the devices that take
3 advantage of DMA and the memory systems of those computers. More specifically, the
4 problem faced by current DMA modules is the ability to adequately service the growing
5 number of high speed devices as well as their varying data requirements.

6 High performance memory systems preferably provide high bandwidth and prefer
7 large data requests. This is in direct contrast to many devices, which may request small
8 amounts of data, have low bandwidth, and require small latencies. This results in system
9 inefficiencies as traditional devices individually communicate with the memory system in
10 an effort to bridge this gap. It is possible that many different devices may be
11 simultaneously making small data requests to a memory system that prefers to handle large
12 memory requests. As a result, the performance of the memory system is decreased.

13 This situation makes it difficult for low bandwidth devices, which may have high
14 priority, to effectively interact with high bandwidth devices that may have lower priority.
15 For example, an audio device may support several different channels that receive data from
16 memory. The audio device typically makes a data request to memory for data every few
17 microseconds for those channels. Because devices such as audio devices recognize that
18 they may experience significant latency from the memory system before their request is
19 serviced, the audio device may implement an excessively large buffer to account for that
20 latency.

21 This is not an optimum solution for several reasons. For instance, many devices
22 maintain a large buffer because they do not have a guarantee that their data requests will be
23 serviced within a particular time period. Other devices maintain an excessively large
24 buffer because it is crucial that the data be delivered in a timely manner even though the

1 devices may have low bandwidth requirements. For example, if an audio device does not
2 receive its data in a timely manner, the result is instantly noticed by a user. Additionally,
3 each device must implement DMA control logic, which can be quite complex for some
4 devices. In other words, the DMA control logic is effectively repeated for each device.

5 Current devices often interact with DMA systems independently of the other
6 system devices and each device in the system is able to make a data request to the DMA at
7 any time. As a result, it is difficult to determine which devices need to be serviced first.
8 The arbitration performed by systems employing isochronous arbitration often defines
9 fixed windows in which all devices that may require servicing are given a portion. These
10 fixed windows are large from the perspective of high bandwidth devices and small from
11 the perspective of low bandwidth devices. Thus, high bandwidth devices are required to
12 buffer more data than they really need and low bandwidth devices often do not need to use
13 their allocated portion of the window. This results in inefficiencies because all of the
14 available bandwidth may not be used and additional memory is required for the buffers of
15 high bandwidth devices. In essence, current systems do not adequately allow high priority
16 devices to efficiently coexist with high bandwidth devices.

SUMMARY OF THE INVENTION

The present invention provides a DMA engine that manages the data requirements and requests of system devices. The DMA engine includes a data reservoir that effectively consolidates the separate memory buffers of the devices. In addition to consolidating memory, the DMA engine provides centralized addressing as well. The data reservoir is divided into smaller portions that correspond to each device. The DMA engine also provides a scalable bandwidth and latency to the system devices. An overall feature of the present invention is the ability to guarantee that a particular device will be serviced in a programmable response time. This guarantee enables the buffer sizes to be reduced, which conserves memory, as well as permits the available bandwidth to be efficiently utilized.

Because the DMA engine maintains the data reservoir, the DMA engine is responsible for providing each device with the data that the device requests. At the same time, the DMA engine is also responsible for monitoring the remaining data in the data reservoir such that a data request can be made to the system's memory when more data is required for a particular portion of the data reservoir. To accomplish these tasks, the DMA engine provides arbitration functionality to the devices as well as to the memory.

The arbitration functionality provided to the devices determines which devices are eligible to make a data request in a particular cycle. Each device may have multiple data channels, but the device is treated as a unit from the perspective of the DMA engine. By only allowing some of the devices to be eligible during a particular cycle, all devices are ensured of being serviced within a particular time period and high bandwidth devices are not permitted to consume more bandwidth than they were allocated.

The arbitration functionality provided between the DMA engine and the memory occurs on a per channel basis rather than a per device basis. Each channel is evaluated in

1 turn to determine whether a data request should be made to memory or whether the
2 channel can wait until it is evaluated again in the future. Because the number of channels
3 is known and because the time needed to service a particular channel is known, each
4 channel is assured of being serviced within a particular time period. This guarantee
5 ensures that the data reservoir will have the data required by the system devices.

6 The arbitration interface between the system memory and the DMA engine
7 addresses the data needs of each channel in a successive fashion by using a list that
8 contains at least one entry for each channel. The DMA engine repeatedly cycles through
9 the entries in the list to evaluate the data or memory requirements of each channel. In
10 addition, the order in which the channels are evaluated can be programmed such that high
11 bandwidth devices are serviced more frequently, while low bandwidth devices are serviced
12 within a programmable time period. Thus, data requests to or from memory are for larger
13 blocks of data that can withstand some latency.

14 Additional features and advantages of the invention will be set forth in the
15 description which follows, and in part will be obvious from the description, or may be
16 learned by the practice of the invention. The features and advantages of the invention may
17 be realized and obtained by means of the instruments and combinations particularly
18 pointed out in the appended claims. These and other features of the present invention will
19 become more fully apparent from the following description and appended claims, or may
20 be learned by the practice of the invention as set forth hereinafter.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

DETAILED DESCRIPTION OF THE INVENTION

The present invention relates to systems and methods for servicing and managing the data requests and memory requirements of devices operating within a computer system. A Direct Memory Access (DMA) engine acts as an intermediary between the memory system and the devices by consolidating the buffer requirements of the devices, providing scalable bandwidth and latency to both the devices and the memory system, minimizing the buffering requirements of the devices through guaranteed scheduling, and efficiently using idle time periods.

An overall feature of the DMA engine is the ability to support the data requirements of the devices in a particular system while ensuring sufficient response time and bandwidth for each device. The DMA engine includes a centralized data reservoir or buffer that replaces the buffers of the individual devices. In addition to reducing or eliminating the need for buffers in the various devices, the consolidated data reservoir of the DMA engine also provides centralized addressing. Also, by centralizing the buffer requirements into the data reservoir, the DMA engine is able to implement the DMA control logic a single time, whereas each device previously required separate DMA control logic.

Another feature of the DMA engine is related to the latency that devices often experience when interacting with memory. The DMA engine ensures that a request from a particular device for data will be handled within a pre-determined time period in part by maintaining the data reservoir that holds each device's data. The data reservoir is maintained on a per channel basis by evaluating factors such as the bandwidth requirements of each channel associated with each device, the anticipated response time of the memory system to service the request of each channel, how long the viable data

1 remaining in the data reservoir will last for each channel, and the like. This information is
2 used to determine whether the channel being evaluated should be serviced immediately or
3 whether the channel can wait until it is evaluated again before it is serviced. In this
4 manner, the DMA engine ensures that each device or channel will have sufficient data
5 stored in the data reservoir.

6 The DMA engine further ensures that the data requirements of all devices will be
7 met within a certain time period by providing an interface to the DMA engine for both the
8 devices and the memory. The DMA engine interface with the memory is adapted to the
9 characteristics of a high performance memory system, while the DMA engine interface
10 with the devices is adapted to the requirements of the devices. The DMA engine is
11 therefore capable of accessing relatively large blocks of data from the memory while
12 providing relatively smaller blocks of data to the devices from the data reservoir.
13 Effectively, the DMA engine permits high priority devices, which may have low
14 bandwidth requirements, to efficiently coexist with high bandwidth devices that may have
15 lower priority.

16 The present invention extends to both methods and systems for servicing the
17 memory requirements of multiple devices. The embodiments of the present invention may
18 comprise a special purpose or general purpose computer including various computer
19 hardware, as discussed in greater detail below.

20 Embodiments within the scope of the present invention also include computer-
21 readable media for carrying or having computer-executable instructions or data structures
22 stored thereon. Such computer-readable media can be any available media which can be
23 accessed by a general purpose or special purpose computer. One example of a special
24 purpose computer is a set top box. Exemplary set top boxes include, but are not limited to,

1 analog and digital devices such as satellite receivers, digital recording devices, cable
2 boxes, video game consoles, Internet access boxes, and the like or any combination
3 thereof. By way of example, and not limitation, such computer-readable media can
4 comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk
5 storage or other magnetic storage devices, or any other medium which can be used to carry
6 or store desired program code means in the form of computer-executable instructions or
7 data structures and which can be accessed by a general purpose or special purpose
8 computer. When information is transferred or provided over a network or another
9 communications connection (either hardwired, wireless, or a combination of hardwired or
10 wireless) to a computer, the computer properly views the connection as a computer-
11 readable medium. Thus, any such a connection is properly termed a computer-readable
12 medium. Combinations of the above should also be included within the scope of
13 computer-readable media. Computer-executable instructions comprise, for example,
14 instructions and data which cause a general purpose computer, special purpose computer,
15 or special purpose processing device to perform a certain function or group of functions.

16 Figure 1 and the following discussion are intended to provide a brief, general
17 description of a suitable computing environment in which the invention may be
18 implemented. Although not required, the invention will be described in the general context
19 of computer-executable instructions, such as program modules, being executed by
20 computers in network environments. Generally, program modules include routines,
21 programs, objects, components, data structures, etc. that perform particular tasks or
22 implement particular abstract data types. Computer-executable instructions, associated
23 data structures, and program modules represent examples of the program code means for
24 executing steps of the methods disclosed herein. The particular sequence of such

1 executable instructions or associated data structures represent examples of corresponding
2 acts for implementing the functions described in such steps.

3 Those skilled in the art will appreciate that the invention may be practiced in
4 network computing environments with many types of computer system configurations,
5 including personal computers, hand-held devices, multi-processor systems,
6 microprocessor-based or programmable consumer electronics, network PCs,
7 minicomputers, mainframe computers, and the like. The invention may also be practiced
8 in distributed computing environments where tasks are performed by local and remote
9 processing devices that are linked (either by hardwired links, wireless links, or by a
10 combination of hardwired or wireless links) through a communications network. In a
11 distributed computing environment, program modules may be located in both local and
12 remote memory storage devices.

13 With reference to Figure 1, an exemplary system for implementing the invention
14 includes a general purpose computing device in the form of a conventional computer 20,
15 including a processing unit 21, a system memory 22, and a system bus 23 that couples
16 various system components including the system memory 22 to the processing unit 21.
17 The system bus 23 may be any of several types of bus structures including a memory bus
18 or memory controller, a peripheral bus, and a local bus using any of a variety of bus
19 architectures. The system memory includes read only memory (ROM) 24 and random
20 access memory (RAM) 25. A basic input/output system (BIOS) 26, containing the basic
21 routines that help transfer information between elements within the computer 20, such as
22 during start-up, may be stored in ROM 24.

23 The computer 20 may also include a magnetic hard disk drive 27 for reading from
24 and writing to a magnetic hard disk 39, a magnetic disk drive 28 for reading from or

1 writing to a removable magnetic disk 29, and an optical disk drive 30 for reading from or
2 writing to removable optical disk 31 such as a CD-ROM or other optical media. The
3 magnetic hard disk drive 27, magnetic disk drive 28, and optical disk drive 30 are
4 connected to the system bus 23 by a hard disk drive interface 32, a magnetic disk drive-
5 interface 33, and an optical drive interface 34, respectively. The drives and their
6 associated computer-readable media provide nonvolatile storage of computer-executable
7 instructions, data structures, program modules and other data for the computer 20.
8 Although the exemplary environment described herein employs a magnetic hard disk 39, a
9 removable magnetic disk 29 and a removable optical disk 31, other types of computer
10 readable media for storing data can be used, including magnetic cassettes, flash memory
11 cards, digital video disks, Bernoulli cartridges, RAMs, ROMs, and the like.

12 Program code means comprising one or more program modules may be stored on
13 the hard disk 39, magnetic disk 29, optical disk 31, ROM 24 or RAM 25, including an
14 operating system 35, one or more application programs 36, other program modules 37, and
15 program data 38. A user may enter commands and information into the computer 20
16 through keyboard 40, pointing device 42, or other input devices (not shown), such as a
17 microphone, joy stick, game pad, satellite dish, scanner, or the like. These and other input
18 devices are often connected to the processing unit 21 through a serial port interface 46
19 coupled to system bus 23. Alternatively, the input devices may be connected by other
20 interfaces, such as a parallel port, a game port or a universal serial bus (USB). A monitor
21 47 or another display device is also connected to system bus 23 via an interface, such as
22 video adapter 48. In addition to the monitor, personal computers typically include other
23 peripheral output devices (not shown), such as speakers and printers.

1 The computer 20 may operate in a networked environment using logical
2 connections to one or more remote computers, such as remote computers 49a and 49b.
3 Remote computers 49a and 49b may each be another personal computer, a server, a router,
4 a network PC, a peer device or other common network node, and typically include many or
5 all of the elements described above relative to the computer 20, although only memory
6 storage devices 50a and 50b and their associated application programs 36a and 36b have
7 been illustrated in Figure 1. The logical connections depicted in Figure 1 include a local
8 area network (LAN) 51 and a wide area network (WAN) 52 that are presented here by way
9 of example and not limitation. Such networking environments are commonplace in office-
10 wide or enterprise-wide computer networks, intranets and the Internet.

11 When used in a LAN networking environment, the computer 20 is connected to the
12 local network 51 through a network interface or adapter 53. When used in a WAN
13 networking environment, the computer 20 may include a modem 54, a wireless link, or
14 other means for establishing communications over the wide area network 52, such as the
15 Internet. The modem 54, which may be internal or external, is connected to the system bus
16 23 via the serial port interface 46. In a networked environment, program modules depicted
17 relative to the computer 20, or portions thereof, may be stored in the remote memory
18 storage device. It will be appreciated that the network connections shown are exemplary
19 and other means of establishing communications over wide area network 52 may be used.

20 As used herein, "data request" refers to either a read or a write operation. Data
21 request can also apply to the interaction between the DMA engine and the system devices
22 or to the interaction between the DMA engine and the main memory of the system. The
23 present invention is primarily discussed in terms of memory reads, but it is understood to
24 apply to memory writes as well. The memory or data requirements of a particular device

1 can be evaluated from the perspective of either the DMA engine or the main memory of a
2 system.

3 Figure 2 is a block diagram that illustrates a DMA engine for servicing and
4 managing the memory or data requirements of system devices. Each device can be a
5 hardware device or a software module or a combination thereof. The devices 220 interface
6 with the DMA engine 200 through a devices interface 250. The devices interface 250
7 allows the DMA engine 200 to service the data requirements of the devices 220 while
8 providing sufficient response time and bandwidth for the devices 220. The devices
9 interface 250 further provides arbitration functionality to the devices 220 such that the
10 DMA engine 200 services the data requests of eligible devices included in the devices 220
11 in any given cycle. In other words, the devices interface 250 determines which devices are
12 eligible to make a service request to the DMA engine 200 in a given cycle or window. In
13 this context, the data requests refer to reading or writing data to the DMA engine 200.

14 As described, the devices interface 250 makes a determination as to eligibility on a
15 per device basis and does not consider the channels that may be associated with each
16 device. The memory interface 270, however, determines whether to make a data request to
17 memory 210 on a per channel basis. The memory interface 270 determines whether a
18 particular channel should be serviced and provides arbitration functionality between the
19 DMA engine 200 and the memory 210. The memory channel evaluates each channel in a
20 repetitive fashion. In this manner, each channel is effectively guaranteed to be serviced
21 within a particular time period. In this context, a data request refers to the transfer of data
22 from the main memory to the DMA engine or from the DMA engine to the main memory.
23 Thus, when a device makes a data request, it does not imply that data is transferred to or
24 from the main memory. Also, when a data request is serviced by the main memory, it does

1 not imply that a device has received or transferred data to the DMA engine even though
2 these actions can occur at the same time.

3 In one example, the memory interface 270 may be viewed as a state machine that
4 produces an output for a given input. The output is whether the channel being evaluated
5 should be serviced and the input includes factors that determine whether the channel is
6 critical. Those factors include, but are not limited to, the amount of data currently
7 available to the channel in the DMA engine, how long it takes the main memory to service
8 the data request of the channel, how long before the channel is evaluated again, and the
9 like. After one channel has been evaluated, the state machine advances to the next
10 channel.

11 After a particular sequence of channels has been evaluated, the state machine
12 begins the evaluation process again at the beginning of the sequence. It is possible for a
13 sequence to include a single channel more than once. While the devices interface 250 and
14 the memory interface 270 are illustrated as being separate from the DMA engine 200, it is
15 understood that the devices interface 250 and the memory interface 270 may be integral
16 modules of the DMA engine 200. In addition, the devices interface 250 and the memory
17 interface may comprise both hardware and software components.

18 Figure 3 is a more detailed diagram illustrating the interaction between the devices
19 220, the memory 210 and the DMA engine 200. The exemplary system illustrated in
20 Figure 3 has devices 220 including device 221, device 222, device 223, and device 224. It
21 is understood that the actual number of devices in a particular system is not limited to the
22 illustrated devices but can vary depending on the configuration of the system. Each of the
23 devices 221, 222, 223, and 224 has one or more channels over which data may be
24

Whenever a device included in the devices 220 requires service for any of the channels of the device, a data request is sent to the DMA engine 200 through the device interface 250. The device interface 250, rather than performing arbitration on a per

1 channel basis, arbitrates the data requests it receives on a per device or unit basis. If a
2 device needs to make a data request for more than one channel, the device is responsible
3 for making a data request for the higher priority channel because a device can usually only
4 make a single request. From the perspective of the DMA engine 200, the bandwidth
5 requirement of each device is determined by the device's channels, and the DMA engine
6 200 uses the latency of the urgent channel as the device latency when considering the
7 device request.

8 The device interface 250 provides arbitration functionality that determines which
9 devices or data requests are eligible to be serviced by the DMA engine 200. Once the
10 eligible devices are identified, a basic arbitration scheme may be used to determine which
11 data request should be granted. Determining which devices are eligible, however, includes
12 scheduling the devices such that latencies can be effectively guaranteed. In addition,
13 scheduling the devices in this manner prevents a particular device from consuming the
14 available bandwidth until other devices have been serviced. Scheduling the devices will be
15 discussed further with reference to Figure 6.

16 In essence, the devices interface 250 provides a calculated latency and bandwidth
17 tradeoff. A device having both a high priority and a low bandwidth may be able to
18 withstand a larger latency than a device having a lower priority and a higher bandwidth.
19 Proper scheduling ensures that high priority devices will have an adjustable, guaranteed
20 response time while reducing the buffering requirements for the high bandwidth device.
21 For example, audio devices are typically considered to be high priority devices and an
22 MPEG device is a low priority device with high bandwidth. Because the MPEG device
23 will be serviced in a programmable response time, the buffer requirement of the MPEG
24 device is reduced even though other devices have to be serviced. A key aspect of the

1 devices interface 250 is that each device is guaranteed of being serviced in a defined and
2 programmable response time.

3 The devices are preferably managed by the DMA engine on a per device basis
4 rather than a per channel basis because many of the devices may have low bandwidth and
5 it is more efficient to consider the bandwidth of all the channels of a device. The memory
6 interface 270, however, uses a list structure to manage the memory or data requirements of
7 the individual channels. The entries in the list structure are channel identifiers that identify
8 the channels of the devices 220.

9 The list, which is described in more detail with reference to Figure 4, may be
10 viewed as a circular list that is advanced to the next entry each time an entry or channel has
11 been evaluated or serviced. Each channel represented by an entry in the list is evaluated
12 for service on a regular basis, and each channel is assured of being serviced in a
13 programmable response time. One reason the response time is programmable is because
14 each channel can be included in the list structure more than once. This enables those
15 channels that need more frequent servicing to be accommodated while still ensuring that
16 the other channels will be evaluated or serviced within a known response time.

17 The DMA engine 200 uses the data reservoir 202 as a memory buffer for the
18 devices 220. As the memory interface 270 rotates through the circular list maintained by
19 the memory interface 270 and evaluates the channels represented by the entries in the
20 circular list, the data remaining in the data reservoir 202 for each channel is evaluated.
21 More specifically, the DMA engine 200 evaluates the portion of the data reservoir 202 that
22 corresponds to the channel in the circular list of the memory interface 270 that is being
23 examined.

1 The criteria for evaluating each portion of the data reservoir 202 include, but are
2 not limited to, how many bytes are left in the portion of the data reservoir 202, a buffer
3 time that corresponds to the rate at which the remaining data is being used by the device as
4 well as how long those bytes will last, the latency of the memory system experienced while
5 accessing the data from the memory 210, and an entry time representing when will the
6 channel be evaluated again. These factors determine whether the channel being examined
7 is critical or requires service. If the channel requires service a data request is made to the
8 main memory. If the channel is not critical, then the channel can wait until it is evaluated
9 again by the memory interface of the DMA engine. One benefit of examining each
10 channel independently of the other channels is that the data can be managed in memory
11 rather than in registers, which results in improved performance.

12 Figure 4 is a block diagram that represents the arbitration functionality between the
13 DMA engine and the memory that is provided by the memory interface 270, which is
14 included in the DMA engine 200. Figure 4 illustrates the memory interface 270, which
15 includes in this example, a main list 271 and a sub list 272. Each entry in the main list 271
16 corresponds to a channel. In a previous example, the DMA engine supported 56 channels,
17 which are represented in the main list as entries or channel identifiers having the values of
18 0 to 55. The channel identifiers are represented as channels 273, 275, 276, 277, 278, and
19 279. It is understood that the length of the main list 271 can vary and only a few entries
20 are illustrated in Figure 4. Each channel identifier can be listed multiple times on the main
21 list 271, but it is preferable that multiple entries for a single channel be evenly spaced on
22 the main list 271. This allows a wide range of programmed response times to be
23 implemented without requiring significant storage or memory. Also, this ensures that the
24 entry time or the time until the channel is to be evaluated again is known.

1 The main list 271 also supports identifier numbers higher than the number of
2 channels supported by the DMA engine. In this example, 8 additional channel identifiers
3 are supported and are represented by the numbers 56 through 63. Seven of these channel
4 identifiers indicate a jump or a call from the main list 271 to a sub list such as the sub list
5 272. The sub-list call 274 is an example of these identifiers and sub-list call 274 points to
6 the sub list 272. The sub list 272 contains channel entries similar to the entries on the main
7 list 271, and each time a call to the sub-list is made, one entry in the sub-list is evaluated.
8 After one entry on the sub-list has been serviced, the next entry in the main list 271 is
9 evaluated and serviced as indicated by arrow 290. The next time a call to the sub-list is
10 made from the main list 271, the successive entry in the sub list 272 is performed.

11 This provides the significant advantage of using smaller tables to replace a single
12 larger table. In Figure 5, for example, if a main list 271 had channels M0, M1 and M2 and
13 the sub-list 272 had channels S0, S1, S2, S3, and S4, then the calling order of the entries in
14 both lists would be M0, M1, M2, S0, M0, M1, M2, S1, M0, M1, M2, S2, M0, M1, M2, S3,
15 M0, M1, M2, and S4. If a single list were used to implement this example, 20 entries
16 would be needed in the list. By using a main list and a sub-list, however, only nine entries
17 are needed in this example: a four entry main list and a five entry sub-list.

18 As illustrated in the previous example, only one entry on the sub-list is evaluated
19 on the sub-list each time a call is made to that sub-list. Thus, another significant advantage
20 of the sub list 272 is that the sub list 272 may be used to hold channels that can withstand
21 longer latencies. Another advantage of the sub list 272 is that the main list 271 may be
22 significantly shorter when sub lists are employed. Otherwise, the main list 271 would have
23 to contain space for the entries on the sub list each time a jump to the sub list occurs.
24 Thus, the use of sub lists conserves memory.

1 With reference to both Figures 3 and 4, assume that channel 273 is an identifier for
2 one of the channels of the device 221. Also assume that the DMA engine 200 maintains
3 the device buffer 204 for the channel 273. When the main list 271 reaches the channel
4 273, the channel 273 is evaluated to determine whether a data request should be made to
5 memory 210. In the evaluation, the channel 273 is first checked to determine basic
6 information such as whether the channel is enabled and which way the data is flowing,
7 either to or from the memory 210. Next, full configuration data of the channel 273 is
8 accessed from a memory channel control to determine the bandwidth requirement, the time
9 until the channel 273 will next have an opportunity for service, the data format, the access
10 style, and the like.

11 Next, the available data for the channel in the device buffer 204 is determined by
12 accessing, for example, memory pointers. The amount of available data, in conjunction
13 with how fast the available data is being used, determines how much time is represented by
14 the available data. This value is compared against the response time, which includes how
15 long until the channel will next be examined, as well as an allowance for system overhead.
16 If the comparison indicates that the time remaining to the channel 273 is less than the
17 response time, then the channel 273 is considered critical and a data request for service is
18 posted by the DMA engine 200. If the channel 273 is critical, the data request is placed in
19 a critical request queue for servicing. If the channel 273 is not critical, the data request
20 may be placed in a non-critical request queue.

21 The critical requests, which are stored in the critical request queue, are then
22 processed or serviced. The critical request queue is preferably a first in first out (FIFO)
23 queue that may be reordered on occasion. In one example, the first four data requests in
24 the queue are examined and serviced in an optimal order. The critical queue stores, in this

1 example, the channel identifier; and control information including, but not limited to,
2 current memory page address, first memory sub-page address, current memory length,
3 transaction size, data format, data access style, and the like.

4 The non-critical request queue is not essential to the operation of the invention, but
5 is used to hold the most pressing non-critical data requests. This queue is able to improve
6 memory efficiency by making use of available cycles. For example, if the critical request
7 queue is empty, then data requests in the non-critical queue may be serviced. Data
8 requests in the non-critical queue may remain indefinitely if there is a large volume of
9 other system traffic. If a request in the non-critical queue becomes critical, it is moved to
10 the critical queue for servicing.

11 When determining the response time for a particular channel, it is often necessary
12 to compute the worst case scenario for that channel. This is often dependent on several
13 factors, including, but not limited to the response time of the memory system, the
14 transaction size and the like. In order to determine whether a particular channel should be
15 serviced involves an analysis of several factors, including but not limited to, the time until
16 the channel will be checked again, the number of requests in the critical queue before a
17 request is posted, the worst case latency from when a requests is posted until it is granted
18 by the memory; and the worst case latency from when a request is granted until its
19 servicing is complete. Some of these factors are design constants while others are
20 dependent on the channel.

21 Because the main list 271 is embodied as a circular list, and because the worst case
22 situations are considered, it is possible to guarantee that a particular channel will be
23 serviced within a certain time period or frame. The advantage of this system is that the
24 data requests to memory from the DMA engine are more suited to the characteristics of the

1 high performance memory. Thus, the DMA engine preferably makes larges requests,
2 accommodates large bandwidth, and is capable of experiencing significant latency without
3 having an impact on the devices.

4 Figure 6 illustrates the arbitration functionality provided by the device interface
5 250. Figure 6 illustrates device 221, which has channels 301, 302, and 303, and device
6 222, which has channels 304, 305, and 306. In this example, the DMA engine 200 requires
7 that the device 221 send a data request 307 whenever any one of the channels 301, 302, or
8 303 of the device 221 needs servicing. Similarly, the device 222 sends a data request 308
9 whenever one or more of the channels 304, 305 or 306 requires servicing. Because a
10 request can represent one of several channels, the arbitration performed by the devices
11 interface 250 is per device rather than per channel. Each device therefore has the
12 responsibility of indicating which channel is most urgent or critical, and the latency that
13 the device can experience is determined from the urgent channel.

14 The device interface 250 has an arbitration mechanism that is used to determine
15 which devices are eligible to make requests to the DMA engine 200. In other words, a data
16 request can only be made to the DMA engine when a device is eligible to make a request.
17 In this example, the arbitration mechanism includes an arbitration count 251 that is
18 represented by four bits, but other representations are equally valid. Eligible devices are
19 determined, for example, by the following comparison logic: ((arbitration count XOR
20 devVal) & devMask), where devVal is the device value and devMask is a defined value.

21 Whenever this logic comparison is true for a particular device, that device is
22 eligible to make a data request for data from the data reservoir of the DMA engine. Using
23 this comparison logic, the eligibility of a particular device can be programmed. More
24 specifically, a particular device can be eligible to make a request every cycle, every other

1 cycle, every fourth cycle, every eighth cycle or every sixteenth cycle. This logic also
2 allows the eligibility of the devices to be staggered or scheduled such that any one device
3 does not consume the available bandwidth. As used herein, "cycle" can refer to a defined
4 time window, a certain number of clock cycles, or any other period in which data requests
5 from eligible devices can be made or serviced.

6 For example, the device 221 may only be eligible every time the two least
7 significant bits of the arbitration count 251 are zero. In this situation, the device 221
8 would be an eligible device for only one out of four cycles or arbitration counts. In a
9 similar situation, the device 222 may only be eligible to make a data request every time the
10 two least significant bits of the arbitration count 251 are both ones. In this situation, the
11 device 222 is only eligible for one out of every four cycles. Even though the device 221
12 and the device 222 are only eligible for one out of every four cycles, they are eligible to
13 make a data request on different cycles. In this manner, the requests of the devices can be
14 scheduled in an efficient manner.

15 The present invention may be embodied in other specific forms without departing
16 from its spirit or essential characteristics. The described embodiments are to be considered
17 in all respects only as illustrative and not restrictive. The scope of the invention is,
18 therefore, indicated by the appended claims rather than by the foregoing description. All
19 changes which come within the meaning and range of equivalency of the claims are to be
20 embraced within their scope.

21 What is claimed and desired to be secured by United States Letters Patent is:
22
23
24

1 1. In a system having a data reservoir storing channel data for one or more
2 channels associated with one or more devices, a method for arbitrating data requests of the
3 one or more channels to a main memory, the method comprising the acts of:

4 defining a circular list having a plurality of entries, wherein the entries
5 correspond to the one or more channels;

6 evaluating a channel associated with an entry in the circular list to
7 determine whether the channel requires service; and

8 servicing the channel by requesting channel data from the main memory to
9 replenish the data reservoir if the channel requires service.

10
11 2. A method as defined in claim 1, wherein at least one of the plurality of
12 entries is a call to a sub list having sub entries that correspond to the one or more channels,
13 wherein one of the sub entries is serviced before returning to the circular list.

14
15 3. A method as defined in claim 2, further comprising the act of evaluating
16 each sub entry as calls to the sub list are made from the circular list.

17
18 4. A method as defined in claim 1, wherein the act of evaluating a channel
19 associated with each entry further comprises the acts of:

20 determining an entry time, the entry time representing how long until the
21 channel will be evaluated again;

22 determining a latency, the latency representing how long the main memory
23 will take to respond to a data request;

24

1 determining a buffer time, the buffer time representing how long data stored
2 in the data reservoir maintained by a direct memory access (DMA) engine will last;
3 and

4 determining that the channel requires service if the buffer time is less than
5 the entry time and the latency.

6
7 5. A method as defined in claim 1, further comprising the act of servicing each
8 channel represented by the entries and the sub entries in a programmable response time.

9
10 6. A method as defined in claim 1, wherein the act of servicing further
11 comprises the act of transferring data from the main memory to the data reservoir.

12
13 7. A computer-readable medium having computer-executable instructions for
14 performing the acts recited in claim 1.

1 8. In a system having a main memory storing data for one or more devices, a
2 method for servicing memory requirements of the one or more devices, the method
3 comprising the acts of:

4 generating a data reservoir by a direct access memory (DMA) engine for the
5 one or more devices, the data reservoir maintaining a buffer for each of the one or
6 more devices, wherein the DMA engine communicates with the main memory and
7 the one or more devices communicate with the DMA engine;

8 determining, by the DMA engine, whether a data request should be made to
9 the main memory for each of the one or more devices;

10 for each of the one or more devices requiring service for the data request,
11 requesting from the main memory, by the DMA engine, additional data to replenish
12 each buffer in the data reservoir for each of the one or more devices; and

13 providing each of the one or more devices with access to each respective
14 buffer in the data reservoir.

15
16 9. A method as defined in claim 8, wherein the act of generating a data
17 reservoir further comprises the act of maintaining a channel buffer for each channel
18 associated with each of the one or more devices.

19
20 10. A method as defined in claim 8, wherein the act of determining further
21 comprises the acts of:

22 evaluating an entry in a list, the entry corresponding to a channel of a
23 device, to determine if the channel is critical based on the factors of:
24

1 a buffer time representing how long until the channel buffer in the
2 data reservoir is empty;

3 an entry time representing how long until the channel corresponding
4 to the entry will be evaluated again; and

5 a latency representing a main memory response time;
6 making the data request for the channel if the channel is critical; and
7 refraining from making the data request for the channel if the channel is not
8 critical.

9
10 11. A method as defined in claim 10, further comprising the act of evaluating a
11 next entry in the list, wherein the next entry corresponds to another channel.

12
13 12. A method as defined in claim 10, wherein the act of making the data request
14 further comprises the act of placing the data request in a critical queue.

15
16 13. A method as defined in claim 10, wherein the act of refraining from making
17 the data request further comprises the act of placing the data request in a non-critical
18 queue.

19
20 14. A method as defined in claim 8, wherein the act of providing each of the
21 one or more devices with access further comprises the acts of:

22 maintaining an arbitration count;

23 determining which of the one or more devices are eligible to make a request
24 for data to the DMA engine using the arbitration count; and

1 increasing the arbitration count after all eligible devices have had the
2 opportunity to make the request for data to the DMA engine.

3
4 15. A method as defined in claim 8, further comprising the act of servicing each
5 of the one or more devices in a programmable time period.

6
7 16. A method as defined in claim 8, further comprising the act of replenishing
8 each buffer in the data reservoir within a programmable time period.

9
10 17. A computer-readable medium having computer-executable instructions for
11 performing the acts recited in claim 8.

1 18. In a computer system having a main memory storing data for one or more
2 devices operating within the computer system, a method for making a data request to the
3 main memory on behalf of one or more devices, the method comprising the acts of:

4 generating a main list, wherein the main list has at least one entry and each
5 at least one entry comprises a particular channel of the one or more devices;

6 for each at least one entry, determining if an associated channel is critical by
7 evaluating:

8 a buffer time representing how long a channel buffer associated with
9 the associated channel will last;

10 an entry time representing how long until the associated channel will
11 be evaluated again; and

12 a latency representing a main memory response time, wherein the
13 associated channel is critical if the buffer time is less than the entry time
14 and the latency;

15 placing the data request in a critical queue if the associated channel is
16 critical; and

17 placing the data request in a non-critical queue if the associated channel is
18 not critical.

19
20 19. A method as defined in claim 18, wherein the associated channel
21 corresponds to more than one entry on the main list.

1 20. A method as defined in claim 18, wherein one of the at least one entries
2 comprises a call to a sub list, the sub list returning to the main list after a single entry on
3 the sub list has been serviced.

4
5 21. A method as defined in claim 18, wherein the act of determining if an
6 associated channel is critical further comprises the act of repeatedly evaluating each entry
7 in the main list.

8
9 22. A method as defined in claim 18, further comprising the act of transferring
10 data from the main memory to a data reservoir maintained by a direct memory access
11 (DMA) engine.

12
13 23. A method as defined in claim 18, further comprising the act of transferring
14 data from a data reservoir maintained by a DMA engine to the main memory.

15
16 24. A method as defined in claim 18, further comprising the act of servicing the
17 data request in the critical queue.

18
19 25. A method as defined in claim 18, further comprising the act of evaluating
20 each at least one entry in a programmable response time.

21
22 26. A computer-readable medium having computer-executable instructions for
23 performing the acts recited in claim 18.

24

1 27. In a system including a main memory storing data for one or more devices,
2 a method for arbitrating data requests from the one or more devices, the method
3 comprising the acts of:

4 creating an arbitration mechanism at a direct memory access (DMA)
5 engine;

6 selecting eligible devices from the one or more devices using the arbitration
7 mechanism; and

8 allowing the eligible devices to make data requests to the DMA engine for
9 one channel of each of the eligible devices.

10
11 28. A method as defined in claim 27, wherein the arbitration mechanism is a
12 counter.

13
14 29. A method as defined in claim 28, further comprising the act of incrementing
15 the counter after all of the eligible devices have had an opportunity to make data requests
16 to the DMA engine.

17
18 30. A method as defined in claim 27, wherein the act of selecting eligible
19 devices further comprises the act of performing a logic operation using device identifiers
20 and the arbitration mechanism for each of the one or more devices, wherein each of the one
21 or more devices is an eligible device when the logic operation is true.

22
23 31. A method as defined in claim 27, further comprising the act of allowing all
24 of the devices to make data requests within a programmable time period.

36. A computer-readable medium having computer-executable instructions for performing the acts recited in claim 27.

a memory interface operably connected with the DMA module, wherein the memory interface arbitrates reservoir data requests generated by the DMA module for data from the main memory to replenish the data reservoir.

38. A system as defined in claim 37, wherein the data reservoir further comprises a plurality of device buffers, the plurality of device buffers storing data for the one or more devices, wherein each of the plurality of device buffers is associated with one of the one or more devices.

39. A system as defined in claim 38, wherein each of the plurality of device buffers further comprises at least one channel buffer for each channel associated with each of the one or more devices.

Docket No. 14531.73

1 41. A system as defined in claim 37, wherein the DMA engine guarantees that
2 the device data requests of the one or more devices is serviced within a programmable
3 response time.

4
5 42. A system as defined in claim 37, wherein the memory interface further
6 comprises a circular list having a plurality of entries, each entry representing one of the
7 channels of the one or more devices, wherein the channels are evaluated to determine if the
8 channels are critical.

9
10 43. A system as defined in claim 42, wherein the circular list is linked to one or
11 more sub lists, the one or more sub lists having additional entries.

12
13 44. A system as defined in claim 42, wherein the DMA engine makes the
14 reservoir data request for the channels that are critical.

ABSTRACT OF THE DISCLOSURE

Systems and methods for servicing the data and memory requirements of system devices. A DMA engine that includes a data reservoir for reducing or eliminating device buffers is provided that manages and arbitrates the data requests from the system devices. An arbitration unit is provided that only allows eligible devices to make a data request in any given cycle to ensure that all devices will be serviced within a programmable time period. The data reservoir contains the data buffers for each channel of each device. A memory interface ensures that sufficient data for each channel is present in the data reservoir by making requests to a system's memory based on an analysis of each channel. Analysis factors include how much data is remaining in the data reservoir, how long will that data last, and how long until the channel will be analyzed again. Based on this analysis, a request is either made to the system's main memory, or the channel waits until it is evaluated again in the future. Each channel is thereby guaranteed a response time.

G:\DATA\PAT\WORDPAT\14531.73.doc




FIG. 2

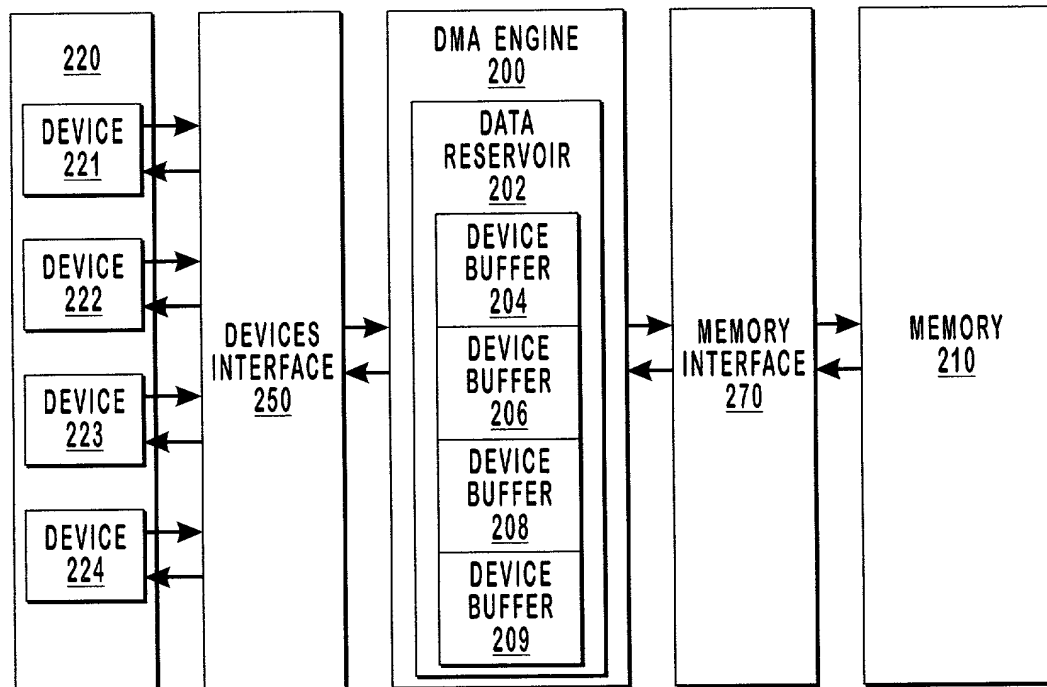


FIG. 3




FIG. 4

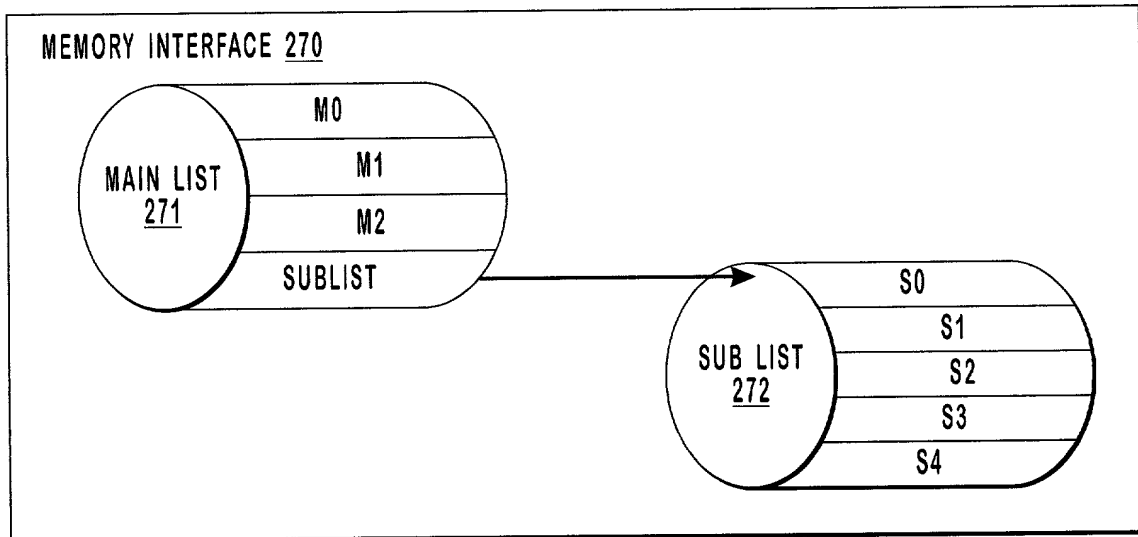


FIG. 5




FIG. 6